BERT, RoBERTa, or DeBERTa? Comparing Performance Across Transformers Models in Political Science Text

Joan C. Timoneda, Purdue University
Sebastián Vallejo Vera, University of Western Ontario

Transformers models such as BERT, RoBERTa, and DeBERTa have revolutionized the field of Natural Language Processing in recent years with substantial improvements in the contextual understanding of text. While political scientists have begun adopting these models, their performance differences are not well understood, especially in cross-lingual applications. This article introduces Transformers models, compares their performance using three different text-as-data political science projects, and shows how to fine-tune them to fit the specific needs of the researcher. We find that RoBERTa and DeBERTa greatly outperform BERT in certain circumstances, and that further training boosts performance in specialized text. In cross-lingual applications, XLM-RoBERTa significantly outperforms both multilingual BERT and multilingual DeBERTa.

One of the most abundant sources of data available to social and political scientists today is text. The rise of social media and open access to Twitter data partly explain this phenomenon, as does recent progress in text digitization techniques that have brought books and manuscripts to digital life that were long hidden away in archives and libraries. Text can help us answer substantive questions in political science on topics as diverse as political campaigns (Hobbs and Lajevardi 2019; McGregor 2020), political polarization and radicalization (Medzihorsky, Littvay, and Jenne 2014), media studies (Matalon et al. 2021), public opinion (González-Bailón and Paltoglou 2015), Supreme Court decisions (Strother 2017), gender and politics (Gleason 2020), and many others. As the availability of text grows, so does the need for computer-based text analysis techniques to supplement those done by humans.

Recent advances in natural language processing (NLP) have spearheaded a text-as-data revolution. In particular, the development of a novel type of deep learning architecture, Transformers, in 2017 has allowed language models such as BERT, RoBERTa, and DeBERTa to understand, classify, and artificially generate text with groundbreaking levels of contextual accuracy (He et al. 2020; Liu et al. 2019; Tunstall, von Werra, and Wolf 2022). Political scientists have begun using Transformers-based models, but there is still a need for clarity around the performance differences across BERT, RoBERTa, and DeBERTa models when applied to political science texts. This is especially true in the case of multilingual classification problems, in which multiple models such as multilingual BERT (mBERT), XLM-RoBERTa, and multilingual DeBERTa (mDeBERTa) exist but there is little consensus over which one performs best or by how much. Lastly, we illustrate how researchers can fine-tune an off-the-shelf Transformers model to apply it to specialized text. First, they can further train the model with new unlabeled text data to suit specific tasks, thus improving contextual understanding of specialized language. Second, they can adapt the model to any specific application with labeled training data. Combining both of these strategies, we argue, yields substantial gains in performance.

In this article, we aim to introduce Transformers models to a broader audience, comparing their performance in English and multilingual models and showing how researchers can further train them on specialized text.[1] Beyond these contributions, we provide detailed evidence for what we believe are

Joan C. Timoneda (timoneda@purdue.edu) is an assistant professor in the Department of Political Science, Purdue University, West Lafayette, IN 47907. Sebastián Vallejo Vera (sebastian.vallejo@uwo.ca) is an assistant professor in the Department of Political Science, University of Western Ontario, London, ON, Canada.

Replication files are available in the *JOP* Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at https://doi.org/10.1086/730737.

1. Further training is the fine-tuning process whereby we provide a model with new words (tokens) and new text on which to learn how the new words are used in context.

Published online January 17, 2025.

the practical advantages of these models for political and social scientists: (1) lower costs to generate data from text through accurate classifiers, (2) potential for large-N analysis for otherwise small-N projects, and (3) the ability to generate new theoretical questions and empirical tests that would not be possible without these models. We highlight their importance through three existing political science projects that use text-as-data for sentence or text classification.[2]

The culmination is three main findings that provide evidence for determining which model to use based on the needs of the researcher. First, RoBERTa is the model that generally offers the greater balance between performance and computational cost. It consistently outperforms BERT by what we consider is a substantial margin. DeBERTa, on the other hand, has similar performance to RoBERTa but uses about twice as much computational power. Second, creating custom-made BERT, RoBERTa, or DeBERTa models through further training yields substantial improvements in classifier accuracy. Third, as expected, all Transformers-based models improve on the performance of other widely used machine learning approaches.[3] With performance gains and the fine-tuning flexibility of Transformers, we believe a growing number of political science projects using text data can benefit from a big data approach.

Indeed, early work applying Transformers-based models showcases the potential of BERT and RoBERTa in political science. Abercrombie et al. (2019) and Abercrombie and Batista-Navarro (2022) employ BERT to detect policy preferences (up to 34 topics) from members of Parliament using debate motions. Similarly, Alemán, Micozzi, and Vallejo Vera (2022) use XLM-RoBERTa to classify legislative speeches by topic. Bonikowski, Luo, and Stuhler (2022) apply RoBERTa to identify frames in US presidential campaigns. Since the focus of these papers is not solely methodological, the selection of their model of choice is unclear (and beyond their scope). Questions also remain around what additional steps scholars could follow to further improve the performance of their models. Our paper provides a systematic overview of Transformers-based models and how to apply them in different classification tasks. While not comprehensively, we show how these models can be fine-tuned for various tasks, including multilingual classification. In comparing the performance of the models, we detail the advantages of using different types of models in examples familiar to social scientists (e.g., performance and computational cost, additional training). Ultimately, we offer an approachable explanation of Transformers models with the aim of making them more accessible to a larger number of applied researchers.

The article is structured as follows. We first briefly introduce the Transformers family and show how it differs from earlier NLP approaches adopted in the social sciences. We then introduce our main arguments in favor of using Transformers-based models in political and social sciences. Three existing text-as-data projects are then used to illustrate our arguments and show the full practical potential of BERT-based models.[4] We conclude with our findings, recommendations, and links to resources for researchers to use these models in their projects.

NLP IN THE SOCIAL SCIENCES

The field of NLP is in the midst of a major revolution. In the past decade, scholars have gone from computing word frequencies and generating broad descriptive assessments to building deep learning systems that understand the contextual meaning of words and sentences (see Grimmer, Roberts, and Stewart 2022; Tunstall et al. 2022). Two factors explain the NLP boom. First, computational power has multiplied recently, providing the necessary technology for computationally intensive text analysis. The second reason is the increased availability of accurate NLP models. Companies like Google and Facebook have invested large amounts of resources to improve text-to-speech and translation technology to detect and weed out (however halfheartedly) certain types of hate speech and disinformation from their platforms. Translation tasks have become increasingly important in an interconnected world. To respond to these needs, new Transformers text models emerged, bringing generational leaps in accuracy and performance.

In political science, machine learning models have become increasingly popular to tackle tasks such as supervised and unsupervised text classification, named entity recognition, sentiment analysis, text similarity, among others. Topic models in particular have been widely used to cluster texts into groups without labeled training data (Catalinac 2016; Grimmer 2010; King, Pan, and Roberts 2013; Roberts, Stewart, and Tingley 2016). These models have worked well with newspaper articles and official statements from political and social elites, but they are less accurate when the text is informal or short (Grimmer et al. 2022). Supervised text classification (the object of this article) uses a labeled training set to train a model that can accurately classify unseen text in the same categories as the training set (Barberá et al. 2021; Pan and Chen 2018).

2. Our goal is in no way to assess or issue any judgments on any of these projects but rather to highlight the alternatives that Transformers-based models offer, especially when compared to traditional human coding approaches and other widely used NLP classification techniques.

3. We compare the performance of Transformers models with SVM and Bi-LSTM RNNs.

4. We will focus solely on supervised sentence classification tasks to make the article tractable, even though the models we introduce have broader applications.

Event extraction from text has also begun using machine learning approaches despite relying on dictionary approaches for years (Beieler et al. 2016; Ward et al. 2013), and named entity recognition has registered improvements in recent years with the growth of libraries such as The Natural Language Toolkit (NLTK) and spaCy.[5]

In supervised text classification, common machine learning approaches include Support Vector Machines (SVM) and Logistic Regression (LR) classifiers. These models use a bag-of-words approach, with an off-the-shelf tokenizer like NLTK to predict the category of a given text.[6] A tokenizer breaks down sentences into tokens that the model can understand, which are then converted into numeric vectors.[7] These vectors enable the model to capture how important words are in a given text sequence. Machine learning models then use these vectorized representations of text to produce classification predictions based on some outcome of interest. While simple, these models can yield good accuracy scores in tasks for which context is not particularly important for performance.

Improving on word vectorization, one of the first revolutionary advances in NLP was Word2Vec word embeddings, developed in 2013 by Google (Mikolov et al. 2017). Word embeddings are mathematical representations of words in a vector space, where vectors closer to each other represent words that are more similar in meaning. A commonly used English-language Word2Vec model trained on a large set of Google News text contains 3 million word embeddings, each of which is a numeric vector of size 1×300.[8] Other widely used word embeddings are Stanford CoreNLP's *GloVe* and Facebook's *fastText* embeddings (Bojanowski et al. 2017; Pennington, Socher, and Manning, 2014).[9] Word2Vec and GloVe embeddings are often paired with recurrent neural network (RNN) and convolutional neural network (CNN) architectures for text classification.[10] One of the most commonly used are long short-term memory (LSTM) networks, a type of RNN that takes in sequential input but keeps important information from further back in the sentence that can be useful to understand new words (Chang and Masterson 2020). Bidirectional LSTM networks, or Bi-LSTMs, process sequential information both forward and backward, further improving contextual understanding. We discuss these alternative models in detail in our first application when we describe the baselines for comparing our main Transformers models.

The problem with these approaches, however, is that word vectors are static, meaning that each word has a corresponding fixed mathematical vector after training. In Word2Vec, GloVE, and fastText, the numeric vector for the word "bear" is the same in "grizzly bear," "teddy bear," "bear fruit," or "bear a loss." This is where the Transformers deep learning architecture innovates: It can dynamically capture the different meaning of words based on context. In the example above, a Transformers-based model would produce four different word embeddings for "bear," one for each specific use of the word. While these approaches are computationally much more expensive, their benefits greatly outweigh the drawbacks. We now introduce these models and their architecture.

THE INTUITION BEHIND TRANSFORMERS NETWORKS

NLP models based on the Transformers deep learning architecture have yielded unparalleled accuracy in sentence classification, question answering, language translation, and other tasks. Transformers are the engines that make BERT, RoBERTa, XLM-RoBERTa, and DeBERTa run.[11] What exactly does the Transformers neural network architecture do to yield these results? They generate dynamic word embeddings that change for every word depending on the context. That is, no word embedding for a given word will be the same across different texts. Each embedding incorporates the word's position in a sentence and its relationship with words that come before and after. The embedding for the word "bear" is different in sentences containing "grizzly bear," "teddy bear," "bear fruit," or "bear a loss." The dynamic nature of these embeddings allows for greater contextual understanding as the context largely determines the word embedding itself. This is the key difference between the Transformers deep learning architecture and other well-known deep learning architectures such as RNNs or CNNs, which mainly use static embeddings.

To accomplish this, Transformers networks use a mechanism called self-attention. This process allows the network to take an entire text as input and process its words all at once

5. The Natural Language Toolkit, or NLTK, is a suite of libraries used to perform many NLP tasks. Text also must be preprocessed first, removing stop words ("the," "and," and so on) and special characters for greater accuracy.

6. More recent models such as Bi-LSTMs use conditional bag-of-words approaches that better preserve context.

7. One of the most popular strategies is NLTK's (https://www.nltk.org), which is often combined with "term frequency–inverse document frequency," or "TF–IDF" vectorizer, for improved performance.

8. As a result, the computational demands of using these large Word2Vec embeddings increased exponentially. However, nowadays, they can be handled easily by most mid- and top-end CPUs.

9. See Rodriguez and Spirling (2022) for embeddings created specifically for political science applications.

10. Neural networks are deep learning algorithms that recognize patterns in data through different layers and nodes. Nodes process information via weights, and each layer produces a new and more accurate representation of the input (see Albawi, Mohammed, and Al-Zawi 2017).

11. Other models, such as XLM-Net and GPT-3, also use the Transformers architecture.

rather than sequentially, as RNNs and previous approaches in NLP do (Tunstall et al. 2022; Vaswani et al. 2017). However, processing all words at once is computationally intensive, and the Transformers architecture is complex. Let us begin by describing its encoder layers. Encoder layers are the different layers in neural networks, including Transformers, that perform computations. Each encoder is identical and includes both the self-attention mechanism and a feed-forward neural network. Information passes from one layer to the next, becoming progressively simpler and more useful by reducing the amount of information and distilling its essence until it can produce the most accurate output for a specific application (say a classification or translation task). Figure 1 describes how the encoder of a Transformers model works. The input ("I love this city") passes through three encoder layers, starting with each token's initial representation or embeddings (x_1 through x_4). These embeddings change as they move from one layer to the next, until the embeddings reach their final form (f_1 to f_4). The diagram in figure 1 shows a three-layer encoder ($encoder_1$ to $encoder_3$), but BERT and RoBERTa base models have 12 layers.

An element within each encoder layer that helps implement the self-attention mechanism is the attention head. Each attention head uses multiple matrices to compute the mathematical relationship between all words in a sentence.[12] In the example of figure 1, the self-attention mechanism allows the model to associate "this" with "city" (rather than with "I"). Thus, the new representation of "this" (z_3 and eventually r_3 in fig. 1) will have some information from the other words in the batch, thus imbuing the embedding of "this" with contextual information from all surrounding tokens. $encoder_1$ will produce one output matrix (composed of r_1 to r_4) that moves on to the following encoder layer ($encoder_2$) and attention head, beginning the process again and repeating it as many times as there are attention heads—12 in total for base BERT and RoBERTa models (see Ravichandiran 2021 and Tunstall et al. 2022). The last encoder outputs a final representation (f_1 to f_4) that can then be decoded to generate a specific outcome such as a translation into another language or a classification into a category (Ravichandiran 2021; Tunstall et al. 2022; Vaswani et al. 2017). The more layers and attention heads, the more precise the final representation.

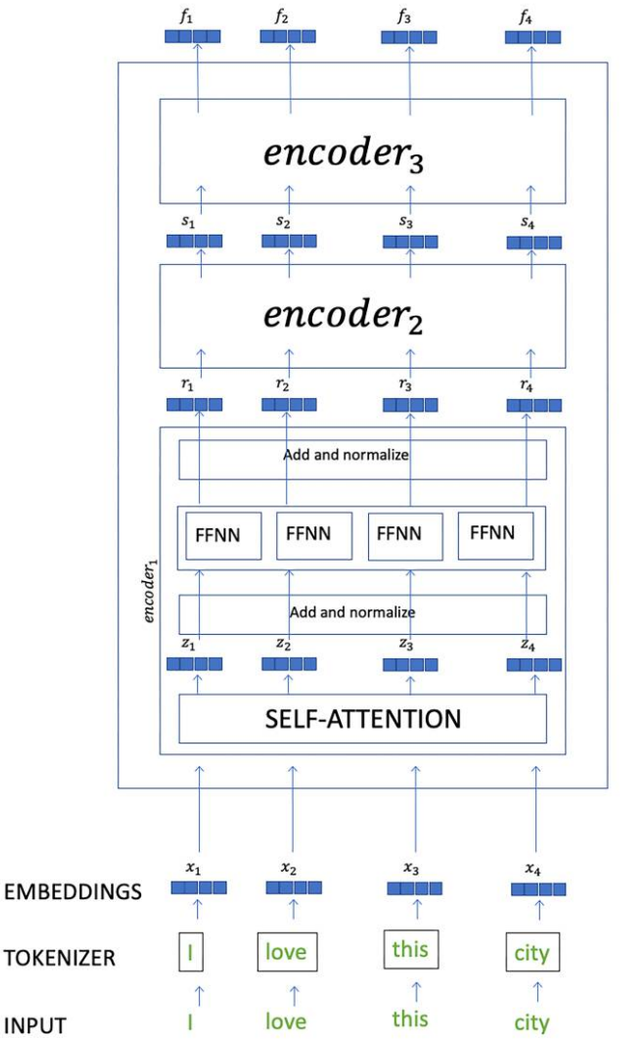

Figure 1. Diagram of a three-stack encoder of a Transformers model. Input text is tokenized and given an initial embedding (vectorized representation) simplified in our figure as x_1 through x_4. The initial embeddings are transformed as they enter the first $encoder_1$. In it, the self-attention mechanism updates the embeddings (z_1 through z_4), which are then passed through a feed-forward neural network. They exit the encoder as a more accurate set of embeddings (r_1 through r_4). The process is repeated for all encoders in the neural network. For example, pre-trained BERT-based models use 12 encoder layers.

As shown in figure 1, the vector representation of each token changes as it progresses through the encoder layers of the Transformers network (Ravichandiran 2021). For instance, the numeric representation for "bear" will be different if it is followed by "fruit" or if it is preceded by "teddy." Each attention head will output an increasingly accurate word embedding for the word "bear" that reflects its meaning in the sentence. In "the child loved the teddy bear because it was soft," the Transformers architecture will use all words in the sentence and give special weight to those it thinks are related to

12. More technically, each input consists of "queries" and "keys" matrices. In the sentence "Mary likes books," to give meaning to the word "Mary" (query), we look at the whole sentence and find the word that is most related to it, in this case, "likes" (key). This process happens for each word in the sentence. Similarly, the self-attention mechanism will estimate the distance (dot product) between the vector representation of every query to the already provided vector representation of every key. Note that, while the tokens can be the same, the representations will be different (as they they come from previous layers). Mathematically, we obtain the dot product of K(eys) and Q(ueries) such that $Attention(Q, K, V) = softmax(QK^T)/\sqrt{d_k})V$, where d is the dimension of K and used as a scaling factor. See appendix F for further details.

"bear"—"child," "teddy," and "soft"—to produce the appropriate embedding for the word "bear," one which captures the idea of a teddy bear rather than a grizzly bear. Since embeddings are numerical representations of words (vectors that follow algebraic rules), the distance between two embeddings is an approximation of the relation between words; for example, vectors that are closer together are more similar than vectors that are farther apart. For example, in the sentence "The kid had to bear the loss of misplacing his teddy bear and his doll," the cosine similarity between (teddy) bear and doll should be higher than the cosine similarity between (teddy) bear and bear (the loss). In appendix H, we provide code and results using RoBERTa to show that this is the case. In appendix F, we provide a more in-depth and technical discussion of Transformers networks.

Lastly, note that we emphasize the difference between Transformers models and Bi-LSTMs in terms of these dynamic word embeddings, but key differences across these models warrant further discussion. First, Transformers process words all at once, using positional embeddings to understand word order.[13] This differs from LSTMs, which have to process tokens sequentially. The second main difference is the self-attention mechanism we described above, which produces contextual word embeddings to fully understand words in context. LSTMs, on the other hand, use static word embeddings. Bi-LSTMs can maintain the meaning of relevant words further back or forward (they are bidirectional) in the sequence to improve predictions, but the embeddings do not change dynamically with context. Because of their capacity to understand context better than traditional LSTMs or CNNs, we expect Bi-LSTMs combined with the most recent state-of-the-art word embeddings (GloVe) to provide the best benchmark on which to compare Transformers models.

THE TRANSFORMERS FAMILY: BERT, RoBERTa, AND DeBERTa

Google AI's BERT and Facebook AI's RoBERTa and XLM-RoBERTa are the encoders of a Transformers model. After all, BERT stands for Bidirectional Encoder Representations from Transformers.[14] Google's BERT encoder consists of 12 encoder layers and 12 attention heads in its BERT-base configuration and 24 encoder layers and 16 attention heads in its BERT-large configuration (Ravichandiran 2021).[15] Also, note the word bidirectional in BERT, which points to its ability to read text forwards and backwards, relating each word to all words in a sentence.

For training, BERT and related models leverage the Transformers architecture, but they also need (1) information and a way to (2) learn from that information. Step 1 is relatively simple, if computationally intensive. The creators of BERT used 11,038 books from the Toronto BookCorpus and all of English Wikipedia to train BERT—a total of 16 GB worth of text (Devlin et al. 2018). Facebook AI's RoBERTa, on the other hand, used the same data as BERT and added more data from Common Crawl (CC-News), Open WebText, and a subset of Common Crawl named Stories,[16] for a total of 160 GB of text, or 10 times more data (Liu et al. 2019). Cross-lingual RoBERTa, XLM-R, was trained using Wikipedia for all languages and data from Common Crawl (Conneau et al. 2019). While it is true that models have increasingly diversified their sources of training data, the over-reliance on Wikipedia and web data has implications for specialized applications, as the language from these websites does not include the types of words and text required for high performance in specialized tasks. Our third application below addresses this issue and provides a solution for researchers to improve task-specific performance.

Step 2 is a bit more complex. How do BERT and similar models learn how words relate to each other in text? Their method for learning is called masked language modeling (MLM).[17] MLM masks about 15% of tokens in a text corpus. Masking means replacing the actual token with <MASK> and then using the full power of the self-attention mechanism from Transformers to predict the masked words. For instance, in the sentence "I love visiting the windy city, <MASK>, the cultural and commercial capital of the Midwest," BERT and RoBERTa will use the information before and after <MASK> to predict "Chicago." While "windy city" may provide a clue, the fact that the city is an important Midwest metropolis is key in predicting the word correctly. MLM takes advantage of the two most innovative and powerful features of Transformers-based models: bidirectionality and self-attention. To predict "Chicago" BERT and RoBERTa use words before and after the <MASK>, not only words before. They use all relevant information in the sentence ("city," "windy," "Midwest," "metropolis," and "commercial capital") to come up with a probability

13. Positional embeddings are an additional matrix where each token is given a number to represent its position within a text.

14. RoBERTa stands for "Robustly optimized BERT approach" and XLM-RoBERTa stands for "Cross-lingual RoBERTa."

15. Each layer of a BERT-based encoder outputs word vectors of length 768, while the BERT-large model outputs word vectors of length 1,024 (the same applies to the base and large versions of RoBERTa and XLM-R). Longer vectors contain a more accurate representation of a word but also require more space and computational power.

16. Common Crawl is a repository of historical websites.

17. BERT also uses next sentence prediction as a training method, but MLM remains the most common training method across the different models.

for the most likely candidate word to replace the <MASK>. During training, BERT and similar models use MLM to predict 15% of all words. Through MLM, these models become highly accurate at word prediction, which means that they can understand all words in a text and how these words relate to one another. In sum, while Transformers is the neural network architecture that produces the most accurate representations through the self-attention mechanism, MLM allows BERT, RoBERTa, and other models to learn about a text.[18] Figure 2 provides a graphical description of MLM.

We highlight four key implications of MLM. First is its ability to understand context well. As explained above, by masking words in text and making the model predict the masked words, we have achieved substantial improvements in performance. The second is the inductive and social biases that emerge from using MLM. Inductive biases help Transformers models using MLM to learn sentence structures, which can lead to downstream performance gains (Zhang and Hashimoto 2021). However, they are also known to "encode worrying levels of social biases," especially around gender and race (Kaneko and Bollegala 2022). Third is the lack of clarity that remains around the share of words in the training data that should be masked. The creators of Transformers models masked 15% of words in the unlabeled training data, which has become common practice. Others have begun to question whether masking more words can yield better results (Wettig et al. 2022). This naturally leads to the fourth implication, namely, that authors require between 0.5 and 1 GB of additional, specialized unlabeled data to make noticeable performance gains.[19] We address this issue further in application 3.

Note also that BERT, RoBERTa, and XLM-RoBERTa all need to first tokenize a sentence before they can compute an output. That is, they need to break the text into words and word chunks that have meaning according to the training process. Since BERT was trained on less data than RoBERTa, it has 30,522 unique vocabulary elements, while RoBERTa has 50,265. The cross-lingual model XLM-RoBERTa, on the other hand, has the largest number of unique vocabulary elements at 250,002 (Conneau et al. 2019). These unique vocabulary elements can handle out-of-vocabulary words by concatenating different word chunks.[20] For instance, the word "training" would be tokenized as "train, ##ing", with the double hashtag indicating that "ing" is a subword token that follows the token "train." This approach to subword tokenization has proven to be accurate in handling out-of-vocabulary words (Ravichandiran 2021; Tunstall et al. 2022).

So far we have discussed BERT, RoBERTa, and XLM-RoBERTa but not DeBERTa. The decoding-enhanced BERT with disentangled attention (DeBERTa) model makes two technical innovations on BERT and RoBERTa (He et al. 2020). One is that it separates the word's content from its position in a sentence and thus computes the final embedding from the Transformers in a way that makes the representation more accurate. The second is that it refines the final decoder layer to achieve better fine-tuning (He et al. 2020).[21] DeBERTa was trained on 78 GB of data and reported improvements between 0.9 and 3.6 percentage points over RoBERTa (large) even when using half of the training data (He et al. 2020). The downside of DeBERTa is that it uses double the graphics processing unit (GPU) RAM of RoBERTa (large) and over three times that of BERT (large).[22] In the Applications section, we compare DeBERTa performance to RoBERTa and BERT to see whether these reported performance benefits outweigh the computational costs.

FINE-TUNING A TRANSFORMERS-BASED MODEL

Where Transformers-based models' abilities shine brightest is in their general applications, a process known as fine-tuning. During fine-tuning, we use a pretrained Transformers-based model and modify its final encoder layer to suit our particular task. Thus, we can harness the great knowledge that BERT and RoBERTa already have about words and text to produce highly

18. BERT models use MLM to estimate the probability distribution for all tokens to replace the <MASK> object. Once we uncover the masked object, BERT can estimate loss, the difference between the probability distributions for each output token and the true labels. In the next iteration, BERT will correct its prediction accordingly. This process appears in more that one stage of the training process. For example, during fine-tuning (explained below), we can use MLM to improve our predictions as well.

19. It could be that by masking a higher percentage of words, MLM would perform better at predicting newly introduced vocabulary elements, as they would be masked more often. Future research should address this issue.

20. Transformers models use different approaches for subword tokenization. BERT uses byte-pair encoding, which reduces words to the character level and creates character groups based on unique vocabulary elements according to their frequency in the training data. RoBERTa uses byte-level pair encoding, which converts characters to bytes (the letter "b" is byte 62, for instance) and then combines the bytes into groups also according to their frequency in the training data. Generally speaking, if a group of characters or bytes exists as a token in the vocabulary, it is used as a token. If not, it is further broken down until a vocabulary token matches it. Because characters are converted to bytes, this method is more helpful in multilingual contexts. These methods differ from techniques such as lemmatization or stemming in that the full word and all the information are kept in subword tokenization, and it is up to the Transformers matrices to calculate how important each subword token is in relation to all other tokens in the sentence.

21. We will not go into depth about the details of DeBERTa's innovations for want of space and an emphasis on clarity and applications. Please refer to He et al. (2020) for more information.

22. This is, to the authors' knowledge, the first article in the discipline to report these findings on computational cost.

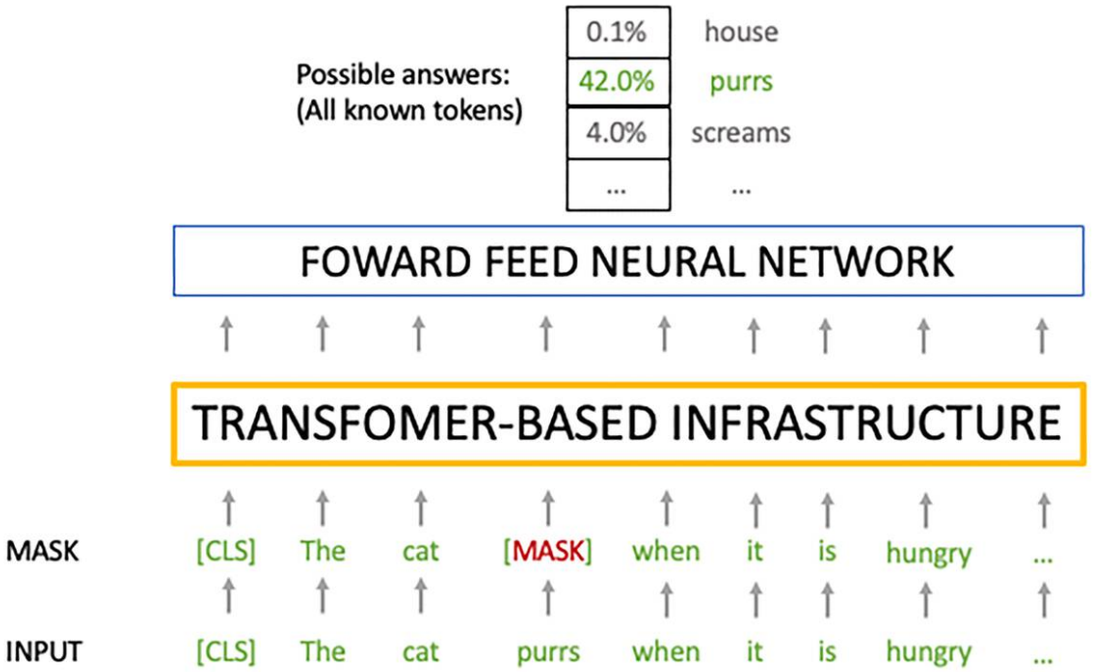

Figure 2. Representation of the masked language modeling. Transformer models randomly mask 15% of tokens. After running the corpus through the Transformers and neural-network architecture, it asks the model to predict the masked word. Transformer models then calculate loss and the required gradient changes to optimize the model's weights and obtain better representations.

accurate sentence classification of campaign slogans, judicial decisions, racism on Twitter, and myriad other important political science research applications. The following subsection details the general procedure to fine-tune and further train Transformers models.

In general, there are two ways to apply these models: standard fine-tuning and further training. Standard fine-tuning is the most common application and refers to using an already trained or off-the-shelf model like BERT and applying it to a specific task using manually labeled training data. Say we are using BERT and trying to classify a text as positive or negative (sentiment). We first need a labeled training set, usually a sample of text manually labeled "positive" or "negative." The size necessary for the training set varies, but one of the advantages of BERT, RoBERTa, and similar models is that they can be fine-tuned using relatively small amounts of training data.[23] We then tokenize all the labeled text using BERT's tokenizer and divide the training data into training and test sets.[24] We set the appropriate model hyperparameters[25] and apply BERT to our specific classification problem. We use 10-fold cross-validation (CV) to fully evaluate the capacity of the model to generalize to unseen data.[26] Once the model has been cross-validated, we train the final version using the full set of training data without splits and obtain our final classified dataset.[27]

The second way to fine-tune a BERT—or similar—model is through further training the model. To do so, researchers provide additional unlabeled text data to the model to improve accuracy for a specialized task. That is, we take all the

23. No hard rule exists on the amount of labeled text per category required to train a Transformers model, but it is ideal to have between 200 and 500 observations depending on the task. Transformers-based models perform well with small training sets (Khan et al. 2021).

24. The training set is a subset of the training data that the model will use to "learn" how to classify the task at hand. The model then uses the test set to evaluate its prediction accuracy on unseen data. Common practice is to split the data into 80% train and 20% test sets in five-fold CV and 90% and 10% in 10-fold CV.

25. Learning rate, number of epochs, batch size, optimizer, and number of warm-up steps.

26. In all our applications, we perform 10-fold CV on the full dataset on a set of 90–10 train–test splits. We do not hold out an extra 10% or 20% of the data for a true out-of-sample test for two reasons. First, we report the test averages for the full CV run as our final test scores, and we do not select a specific "best model" from the CV run. The second, and more constraining, reason is that we do not have large amounts of data for applications 1 and 2, and withholding data from training affects performance noticeably. We further explain CV and how to set hyperparameters in appendix B.

27. Note the difference between the CV and final fine-tuning step. The cross-validated performance metrics give us a measure of the model's performance to generalize to unseen data. The final fine-tuning step uses the full training set without splits for training and test data and the best model hyperparameters from CV. See appendix B for a more detailed discussion on CV and appendix C for a step-by-step guide with Python code snippets on how to perform standard fine-tuning.

knowledge BERT has and add (1) new raw text data and (2) new vocabulary elements to allow it to understand specialized text better. As a practical guide, the researcher must first gain access to a BERT, RoBERTa, or DeBERTa pretrained model that they want to train further. After that, they must add to the model's tokenizer a new set of unique vocabulary elements specific to the domain of study. The choice of vocabulary should be grounded on the researcher's expertise on a topic. Third, the researcher must collect, tokenize,[28] and administer new unlabeled data containing the new vocabulary elements to the base pretrained model. The fourth step is to retrain the model, which allows it to see the new vocabulary elements in context and understand them better. Lastly, the model can then be saved and applied using standard fine-tuning to build a classifier. We provide a more detailed guide with Python code snippets on how to further train a model in appendix E.

For example, in application 3 in the next section, we further train a RoBERTa-large model to better recognize text that contains the words "covid" and "coronavirus," two terms that did not exist in 2018 when Facebook AI trained the original RoBERTa. We use an unlabeled set of 6,076 academic abstracts and 4.8 million tweets and news headlines, all of which are about Covid-19. The idea is that performance improves when adding "covid" and "coronavirus" as new unique vocabulary elements to the model and training it to understand what these two new elements mean in relation to other words in different texts. Note that further training generates a new variant of the BERT or RoBERTa model, which can then be used for standard fine-tuning as described above. We detail how to further train a RoBERTa model in application 3.

Once researchers have built and saved the final classifier, they usually apply it to a new and unseen dataset to generate larger amounts of data. It is important, however, to be cognizant of the potential measurement bias present in Transformers-generated data, as it can lead to estimation bias. We closely follow recommendations by Egami et al. (2023), who propose a method to reduce measurement bias in data generated through large language models. In appendix I, we provide a detailed explanation and illustration of Egami et al.'s (2023) methodology.

In sum, by fine-tuning a Transformers-based model for our own application, we can improve contextual understanding and therefore task-specific performance. We can then apply the custom-built model to unseen data to generate larger amounts of labeled data for analysis, correcting for any potential measurement biases present in the machine-labeled data. We now illustrate the techniques described above using three different applications relevant to political scientists.

28. Using the updated tokenizer with the new vocabulary elements.

APPLICATIONS

We use three projects within or relevant to political science to compare the performance of BERT, RoBERTa, and DeBERTa models in different types of text data. We selected the projects carefully to provide a wide array of potential applications within political science. The first project uses English text from Twitter and produces a binary classification of civil and incivil tweets. The second project uses text in 29 languages and classifies it into four categories. The final project uses Covid-related text in English and a binary classification of true and fake news. This final application helps illustrate the advantages of further training BERT and RoBERTa. We also provide results from two non-Transformers baseline models for reference: (1) SVM, which uses a simpler machine learning model to classify text, and (2) a LSTM RNN, which is the non-Transformers state of the art model in the literature. We provide further details on these two baseline models below.

We run all of our models Python.[29] We use learning rates within the ranges suggested by the authors of RoBERTa (Liu et al. 2019), the weighted Adam optimizer, and 10 warm-up steps.[30] We use four epochs for BERT, five epochs for RoBERTa and DeBERTa, and six epochs for the cross-lingual models, all with an early stopping mechanism to prevent overfitting.[31] The computational requirements for each model vary widely, and researchers should consider the performance–cost trade-off

29. The code uses three common machine learning and deep learning Python libraries: Transformers, Torch, and Scikit-learn. Code for these libraries is widely available and accessible to all applied researchers. We have also made our code available on GitHub (https://github.com/joantimoneda/Bert_Roberta_Deberta_jop). All pretrained models used in this article are publicly available at huggingface.co. These include BERT-large (bert-large-uncased), RoBERTa-large (roberta-large), DeBERTa-large (deberta-v3-large), XLM-RoBERTa (xlm-roberta-large), mDeBERTa (mdeberta-v3-base), and mBERT (bert-base-multilingual-uncased).

30. The learning rate is the parameter that determines how quickly the model "learns." A low learning rate can lead to slow convergence or the model lingering in local optima. A high learning rate can often lead to lack of convergence because the model overshoots the solution. Scholars should monitor loss and accuracy gains to ensure they pick the right learning rate to maximize performance. Recommended learning rates are 3*e*–4, 1*e*–4, 5*e*–5, or 3*e*–5 for BERT-large-uncased (our choice: 3*e*–5); 1*e*–5, 2*e*–5, or 3*e*–5 for RoBERTa-large (our choice: 3*e*–5); and 5*e*–6, 8*e*–6, 59*e*–6, or 1*e*–5 for DeBERTa-V3-large (our choice: 1*e*–5). Each researcher can then move the needle up or down to fit each specific task, testing learning rates in these ranges. For XLM-R, 5*e*–6 is recommended, but again, this may vary slightly between 5*e*–4 and 5*e*–7 in most applications. The batch size is the number of samples that will be propagated through the network in each iteration. A size of 16 is preferred over one of 8 and will produce better performance, especially with more than two labels. After that, a batch size of 32 makes the model run faster but does not meaningfully improve results and is much more computationally demanding (in terms of GPU RAM). Weighted Adam is an adaptive optimizer that helps improve convergence and generalizability (see Loshchilov and Hutter 2017).

31. Training ends after the first increase in validation loss when compared to training loss.

Table 1. Model Performance by Category (Main "Type" Variable; 10× Repeated CV Scores)

	Civil				Incivil				Macro-Average			
Model	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}
ML–SVM	.787	.832	.809	.005	.420	.350	.379	.012	.591	.604	.594	.007
Bi-LSTM–GloVe	.806	.841	.820	.006	.479	.409	.423	.009	.620	.644	.627	.007
BERT-large	.828	.921	.872	.002	.672	.448	.527	.021	.684	.751	.700	.010
RoBERTa-large	.873	.915	.893	.004	.716	.610	.650	.026	.763	.795	.771	.014
DeBERTa-v3-large	.867	.906	.885	.005	.694	.598	.638	.017	.752	.780	.761	.011
Random baseline (F1)							.618					
Majority baseline (F1)							.743					

Note. Random and majority baselines represent macro-average F1 scores.

when deciding on which model and which platform to use.[32] We provide further details on the models in appendix B, including how long each model takes to run and the maximum number of tokens used per application (see table B.1). We also provide the standard deviation for the F1 scores of all the CV runs in the tables 1, 2, and 3.

Lastly, while this article and the applications that follow focus primarily on the predictive performance of Transformers models, they are also useful for measurement. Indeed, we would emphasize the importance of using these models to create new and more accurate measures of complex political phenomena. For instance, Transformers models can help us generate more precise measures for Supreme Court Justices' ideal points, interest group interactions in Latin American parliaments, democratic incumbents' social media use, and a host of other phenomena where text is a primary source of data. In fact, the variables that result from these models, and which are then used in statistical models, are often their most powerful contribution.

Incivility on Twitter

The first project studies incivility in US state legislatures (Gervais and Morris 2019). The authors leveraged increased Twitter activity by state legislatures (the institutions themselves have Twitter accounts) and organized parties in state legislatures between 2006 and 2018, collecting all Twitter activity in this period for all 50 US state legislatures and organized state parties. The process yielded 344,000 total tweets, and the authors built a sample of 2,076 tweets. They used three research assistants to code each tweet as either civil or incivil. Examples of incivil tweets are: ".@AKIndDems Wrong. Please stop lying to the twitterverse. Interested in the truth? Read here- https://t.co/D2p8OZAjBl #akleg" and "@AKIndDems - that's not true. Don't play fast and loose."[33]

The coders agreed on the labeling 75% of the time, and the average agreement across coder pairs was 78.7%. The Kappa score for intercoder reliability is 0.591, indicating moderate agreement. The authors of the article decided to code as "incivil" tweets where at least two coders agree. This led to a total of 534 incivil tweets (25.7%) and 1,542 civil tweets (74.3%) in our final dataset, a slightly imbalanced panel for which we will report F1 scores. The F1 score is the harmonic mean between precision and recall scores and therefore better captures accuracy in unbalanced panels.[34] The mean length for these tweets is 19.29

32. In appendix A, we include a breakdown of the computational requirements and costs associated with each model on table A.1. In this article, we use Jupyter Notebooks that we run either on Google Colab (free GPU acceleration up to 16 GB of GPU RAM) or datacrunch.io for more demanding tasks. Table B.1 in appendix B details which computing platform we use depending on the model and the number of tokens.

33. The authors set four criteria for incivil tweets: (1) name-calling, mockery, sarcasm, and character assassination; (2) spin and misrepresentative exaggeration; (3) emotionality/digital stridency; and (4) conspiracy theory.

34. This is because it considers both how well the model has identified true positives as opposed to generating false positives (precision) and the model's ability to identify true positives as opposed to generating false negatives (recall). Some models may generate a lot of true positives and very few false positives (high precision), but they may also generate a lot of false negatives (low recall) or vice versa. This is a larger problem in unbalanced panels. When only 5% of observations are one, models can produce more false negatives because the zero category dominates, leading to low recall, but they may also have high levels of accuracy because most observations will be classified correctly. By using the F1 score instead of the accuracy score, we get a much more representative picture of the model's performance.

Table 2. Model Performance by Speech Type (10× Repeated 10-fold CV Scores)

	Campaign				Famous				International				Ribbon Cutting				Macro-Average			
Model	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}
ML–SVM	.716	.668	.684	.008	.559	.618	.580	.009	.809	.831	.817	.010	.775	.708	.740	.010	.708	.715	.705	.006
Bi-LSTM– GloVe	.730	.715	.709	.007	.552	.558	.540	.015	.759	.708	.719	.029	.697	.668	.682	.017	.668	.685	.663	.005
mBERT	.794	.777	.783	.002	.722	.668	.689	.004	.844	.870	.854	.011	.819	.853	.832	.002	.792	.795	.789	.002
XLM-RoBERTa	.838	.845	.839	.004	.792	.724	.750	.011	.863	.933	.895	.007	.865	.854	.856	.014	.839	.840	.835	.009
mDeBERTa	.784	.777	.775	.011	.697	.652	.670	.012	.831	.882	.852	.005	.833	.831	.828	.004	.786	.786	.781	.001
Random baseline	.25																			
Majority baseline	.262																			

Note. Random and majority baselines represent macro-average F1 scores.

Table 3. Model Performance on Fake News Dataset, 10× Repeated 10-fold CV

	Fake				True				Undetermined				Macro-Average			
Model	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}	Prec.	Rec.	F1	SD_{F1}
ML–SVM	.717	.749	.731	.010	.701	.731	.713	.006	.743	.680	.708	.007	.720	.720	.717	.004
Bi-LSTM– GloVe	.719	.732	.717	.014	.720	.702	.703	.013	.712	.692	.694	.013	.709	.717	.705	.11
BERT	.805	.766	.782	.004	.745	.806	.772	.009	.748	.714	.724	.003	.762	.767	.760	.003
BERT-Covid	.857	.784	.815	.003	.772	.809	.786	.002	.744	.762	.749	.005	.785	.791	.783	.003
RoBERTa	.872	.778	.819	.010	.753	.827	.783	.009	.748	.745	.742	.011	.783	.791	.781	.008
RoBERTa-Covid	.875	.835	.851	.008	.788	.840	.810	.004	.784	.758	.767	.005	.811	.816	.809	.003
DeBERTa	.865	.820	.838	.007	.754	.832	.788	.008	.798	.738	.761	.006	.797	.806	.796	.004
DeBERTa-Covid	.881	.834	.855	.007	.772	.850	.808	.003	.786	.742	.760	.002	.809	.813	.808	.001
Random baseline									.44							
Majority baseline									.57							

Note. Random and majority baselines represent macro-average F1 scores.

words, resulting in an average of 42.32 tokens per tweet using the RoBERTa tokenizer and a maximum number of tokens of 147. We shuffle the full 2,076-tweet sample and use 10-fold CV to test the ability of the model to generalize to unseen data. We report the average scores from repeated (10 times) 10-fold CV for all models in tables 1 and 2. Repeated 10-fold CV consists of performing 10-fold CV multiple times (in our case 10 times), yielding an average of model performance across 100 models from three different sets of 10 folds. This approach provides a much more accurate estimate of true out-of-sample model performance.[35]

In table 1 we report our first set of tests using the ideal "type" variable where at least two coders agree.[36] We compare the performance of BERT, RoBERTa, and DeBERTa and provide two other non-Transformers baseline models that have been widely used in the literature. First is a machine learning approach using SVM with an NLTK English language tokenizer and a TF-IDF vectorizer.[37] NLTK is a powerful NLP library in Python that helps us convert text into tokens that models can understand (Loper and Bird 2002). The TF-IDF vectorizer is a commonly used tool that penalizes common words and gives particular importance to rare but more meaningful words, which helps the model to understand text better. This model shows how well machine learning models can classify text based on advanced calculations of word frequencies and word importance, but it should perform worse than the rest of the models given its relative simplicity. The second model is a Bi-LSTM RNN, which we pair with GloVe word embeddings (see Chang and Masterson 2020).

The results from table 1 align with our expectations, confirming that Transformers models produce the most accurate classifications for civil and incivil tweets. The F1 score for RoBERTa's classification of civil tweets is high at 0.893. For incivil tweets, on the other hand, RoBERTa's F1 score stands at 0.650 on average over 10-times repeated 10-fold CV (100 models in total). The model's overall F1 accuracy for RoBERTa is 0.771. DeBERTa performs similarly to RoBERTa across all models, but BERT's performance is worse. BERT's F1 score for

35. We use a learning rate of 1e-05 (DeBERTa) and 3e-05 (BERT and RoBERTa), a batch size of 32, and train the models over four, five, or six epochs. Monitoring validation loss versus training loss is a standard approach to determining whether the model is overfitting. When the validation loss exceeds the training loss, the model is trying to increase accuracy on the training data at the expense of generalizability. This is reflected in worse performance on the validation set.

36. This variable tends to produce the best performance across all models.

37. We apply the TF-IDF vectorizer, which weighs token frequencies once document frequencies are also considered, to the SVM models to maximize their performance, as they do not use word embeddings. The TF-IDF vectorizer we use is from the scikit-learn Python library.

incivil tweets is 0.527, a drop-off in performance of 0.123 when compared to RoBERTa's 0.650. This is a statistically significant difference considering both models' SDs. BERT performs well with civil tweets ($F1 = 0.872$), but its overall performance is worse than RoBERTa and DeBERTa.

Compared to the baselines, all Transformers models improve on the performance of simple machine learning and standard neural network approaches. As expected, the Bi-LSTM model with GloVe embeddings performs better than SVMs. Simpler machine learning models cannot understand the linguistic nuances in civil and incivil tweets and produce low levels of accuracy for incivil tweets.[38] RoBERTa still improves on Bi-LSTM-GloVe's performance by 53.2% in the incivil tweets category and 23% in overall model performance (when comparing F1 scores). Note that overall levels of accuracy may appear artificially high for these two models considering that civil tweets are easiest to classify as they represent 74.3% of the sample. Indeed, the random accuracy for the (unbalanced) dataset is 74.3% for zeros (civil) and 25.7% for ones (incivil).[39]

We therefore recommend that researchers fine-tune a RoBERTa model to generate labeled data from a larger set of texts. Equipped with this larger dataset with both human-labeled and machine-labeled data, we then recommend applying the Egami et al. (2023) method in downstream statistical analyses. This helps identify and correct for measurement bias in the machine-labeled data. We provide a full example of this method in application 3 and appendix I.

Classifying multilingual speeches

The second project we use to illustrate the advantages of Transformers-based models of NLP is the Global Populism Database (GPD) introduced by Hawkins et al. (2019).[40] The GPD project started in 2006 with the goal of creating a large dataset of global populist discourse by political leaders. The project currently contains 1,161 speeches by 234 leaders from 73 countries. The speeches are in 29 different languages,[41] and their length ranges from 18 to 20,587 words, with a mean length of 2,449.8 words and a median length of 1,938 words. The languages in the sample are diverse. Pretrained models in high-resource languages—that is, languages with a large amount of data available—such as English, Spanish, and German, use around 100 GB of data. For low-resource languages, such as Albanian, Latvian, and Tagalog, the pretrained data can be closer to 10 GB. The lack of training data for some of these languages poses a general challenge for XLM-RoBERTa models to accurately classify text in certain languages (see Conneau et al. 2019). For us, the challenge is to fine-tune two cross-lingual classifiers using a relatively small sample size (1,161 speeches).

GPD's coders have generated multiple variables from the speech data. We will focus on one of these, the type of speech, which has four categories: international (the audience is foreign and is preferably given outside the country), campaign (usually the opening or closing of the campaign), ribbon cutting (given to a local audience), and famous (a widely circulated speech that shows the leader at his or her best).[42] Classifying the type of speech helps identify the context in which populism most frequently occurs. The distribution of the categories is as follows: 304 speeches are famous (26.2%), 304 are international (26.2%), 294 are ribbon cutting (25.3%), and 259 are campaign (22.3%). A majority of speeches are over 512 words and tokens (99.6%), so we set a maximum length of 512 for the model and select the first 512 tokens.

There are three challenges with this dataset. First, speeches are long (some have over 50,000 words), and Transformers-based models are limited to 512 tokens; with 512 tokens, the computational requirements for graphics acceleration are high. The other problem is sample size.[43] Having 843 speeches across three or four categories leaves only between 200 and 300 speeches per category to train the model and a lower number for the training set after the train–test split required in 10-fold CV.

38. The higher accuracy of SVM and LR in civil tweets is trivial. The models default to predicting civil when they cannot distinguish between the two because it is the category with the most tweets in the sample.

39. Random accuracy refers to the underlying probability that the model classifies an event correctly using only the share of a given category in the data. Random accuracy reflects only a lucky guess by the model without the need for further learning.

40. The full dataset and the codebook can be found at https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/LFTQEZ.

41. Bulgarian, Czech, German, Greek, English, Spanish, Estonian, Finnish, French, Hungarian, Croatian, Italian, Japanese, Lithuanian, Latvian, Macedonian, Dutch, Norwegian, Polish, Portuguese, Romanian, Russian, Slovak, Slovenian, Albanian, Swedish, Tagalog, Turkish, and Ukrainian.

42. See also Lewis, Paul, Sean Clarke, and Caelainn Barr. 2019. "How We Combed Leaders' Speeches to Gauge Populist Rise." *The Guardian*. https://www.theguardian.com/world/2019/mar/06/how-we-combed-leaders-speeches-to-gauge-populist-rise.

43. There are important implications related to the 512-token limit for Transformers models. First, longer texts (say Supreme Court decisions) may have multiple sections, and thus key pieces of information may be scattered. If, say, we set a rule to use the first 512 tokens, we may miss important information for the classifier. This is, in our view, the most relevant drawback of the 512 limit. The authors have found in our testing (both for this article and other independent research) that classifier performance does not improve after using around 300 tokens in texts longer than 512 tokens. This may appear counterintuitive, but in fact, the classifiers are so accurate that it is often the case that around 300 tokens are sufficient to comprehend the meaning of a text, and the increases in performance after that are marginal. The exception are texts in which information may be hidden in sections beyond the 512 limit. For those applications, we recommend preprocessing the text to improve on the rule for which set of tokens to select from the text.

Lastly, this multilingual dataset covers tokens for English, Spanish, Japanese, Chinese, and Albanian, among many others. Considering these challenges, the GPD data provide a strong test for the ability of XLM-R to classify speech types accurately.

For this test, we compare the performance of XLM-R, mBERT, and mDeBERTa. As in the first application, we provide baseline results for SVMs and Bi-LSTM with GloVe embeddings.[44] Note, however, that the SVM's tokenizer (NLTK) and the Bi-LSTM's word embeddings (GloVe) are not cross-lingual, so we have first to translate all speeches into English (we could also translate them into other languages, but English is usually the most accurate). We did this using Google Translate, a service whose main engine is a refined and proprietary mBERT model.[45] We used these models the same way we did with the English text in the first application above. We used all speeches in their original languages for the XLM-R, mBERT, and mDeBERTa models.[46]

Table 2 shows the results. The F1 scores show that all Transformer models substantially improve upon the performance of SVM and Bi-LSTM, which is to be expected. Within the Transformers family, XLM-R performance is particularly impressive. Its F1 score for all types of speeches is much higher than mBERT and mDeBERTa, which score similarly. The largest difference is with famous speeches, which XLM-R classifies 8 percentage points more accurately than mDeBERTa (6.1 when compared with mBERT). This represents an 11.94% increase in performance. XLM-R outperforms the SVM model by 29.3% and the Bi-LSTM by 38.9% in classifying famous speeches. Furthermore, XLM-R outperforms mDeBERTa by around 4.5 percentage points on average in the other three categories. With XLM-R, two of the four speech type categories have F1 scores over 0.839 (campaign and ribbon-cutting speeches) and close to 0.90 (international speeches). High precision and recall scores show that neither false positives nor false negatives are of concern. Across all speech categories, mBERT and mDeBERTa perform better than the non-Transformers models but noticeably worse than XLM-R. This difference is likely because XLM-R is a larger model than mBERT and mDeBERTa, with more training data, tokens, and parameters, making it more accurate for cross-lingual applications with longer texts overall.[47]

These results are particularly encouraging considering the aforementioned data limitations. Even though the classes are well-balanced (around 300 observations in each speech type),[48] there are only 1,161 speeches to train and test the model. Moreover, the speeches are in 29 languages, some of which are trained on much less text data than others. Compounding this, we can only take the first 512 tokens from each speech for classification due to model and GPU constraints. However, the model performs very well, showing the true potential of XLM-RoBERTa and other Transformers-based cross-lingual models in text classification tasks, especially with small training sets.

Given the results above, we recommend that researchers fine-tune an XLM-RoBERTa model over mBERT or mDeBERTa to generate labels in a larger dataset based on the human coded data from the GPD. Again, once the machine-labeled data is available, applied researchers should follow Egami et al. (2023) to detect and correct for any measurement bias in the data before conducting their final analyses (see application 3 and app. I).

Detecting Covid-19 fake news

The third application focuses on detecting fake news around the Covid-19 pandemic and shows the flexibility of Transformers models and how to apply them to specific tasks through further training. This application shows the opportunity that Transformers models hold for increasing performance in specialized domains. While we focus on a specific case here (Covid-19), there are many domains in which researchers can greatly improve classifier performance—and therefore generate better data—by following the steps we outline in this section.[49] We use a manually labeled dataset of true and fake news around the Covid-19 pandemic (see Cheng et al. 2021), an increasingly salient topic of study within political science (Calvo and Ventura 2021; Greer et al. 2021; Timoneda and Vallejo Vera 2021). The authors of the project gathered 7,179 news headlines and Twitter posts containing the words "coronavirus" or "covid" between December 2019

44. We use a learning rate of 5e-06, a batch size of 16, and four epochs for XLM-R and 3e-05, 32 batch size, and four epochs for both mBERT and mDeBERTa. We set the max length to 512 tokens per speech, the maximum we can fit into a Nvidia A100 GPU (80 GB of RAM).

45. We used Google's API and kept the translations to a maximum of 3,000 words per speech to keep the comparison with other models fair and costs down.

46. We do not expect the translation to be a source of bias, considering the improvements in accuracy of machine translation in recent years, especially with English as the target language. It is an inherent weakness of non-Transformers models that they cannot handle text in more than one language at a time.

47. Please see appendix A. For XLM-R, both *xlm–roberta–large* and *xlm–roberta–base*, a smaller model, exist. However, for mBERT and mDeBERTa, only base versions exist. We take the best possible model available from each for this application to show the true power of available cross-lingual models.

48. The exact numbers are as follows: famous, 304; international, 304; ribbon-cutting, 294; and campaign, 259.

49. For instance, researchers have used RoBERTa models to increase performance in the classification of US Supreme Court decisions (Johnson, Strother, and Timoneda 2024) and of racist text in Spanish (Gordillo, Timoneda, and Vallejo Vera 2024).

and September 2020. Through fact-checking websites, they labeled each story as fake, true, or undetermined.[50] Lazer et al. (2018) define "fake news" as "fabricated information that mimics news media content in form but not in organizational process or intent." Examples of fake news from the aforementioned dataset are: "Coronavirus was created in a government lab as a bioweapon and then released on the people of China" or "Japanese doctors advice that taking a few sips of water every 15 mins will prevent the new coronavirus from entering your windpipe and lungs." We use the final dataset, which has 3,681 fake (51.27%), 1,878 true (26.16%), and 1,620 (22.57%) undetermined news stories and tweets. The mean length for the headlines is 21.56 words, and the longest is 143 words, resulting in an average of 29.75 tokens per sentence using the RoBERTa tokenizer and a maximum number of tokens of 160.

Data on the coronavirus pandemic provides a clear example of the advantages of further training Transformers models. When RoBERTa was originally trained in 2018 to 2019, Covid-19 did not yet exist. Coronaviruses had circulated for years, but none had resulted in a global pandemic. The word "Covid" did not exist until February of 2020, so this case provides an intuitive application for how to further train a model on highly specialized words that we can be certain the original model was never trained to understand. Original RoBERTa, therefore, cannot know what Covid-19 is or how the words "coronavirus" or "Covid" are used in context today without further training. Our solution is to further train a new RoBERTa model with new data containing the words "coronavirus" and "Covid" and add those two vocabulary elements to the tokenizer. The new model should be able to classify texts containing these two new words more accurately than original RoBERTa.

There are four steps to further train RoBERTa and other Transformers models. First, we add two new vocabulary elements—"Covid" and "coronavirus"—to the RoBERTa-large (fast) tokenizer, neither of which exists in the original. The tokenizer now has 50,267 unique vocabulary elements, two more than the original's 50,265. Second, we take the pre-existing vector representations for "virus" and "respiratory" from the existing set of RoBERTa vocabulary elements and assign the mean of those two word vectors to be the vector representation for the newly created elements "Covid" and "coronavirus."[51] Third, to make these initial representations more accurate, we feed a set of unlabeled texts containing the words "Covid" and "coronavirus" in English to the model and train it again. In our case, we used 6,079 abstracts from academic articles on the topic of Covid-19 and the coronavirus pandemic. We also added 1 GB of short news headlines around the Covid-19 pandemic obtained from Twitter. The format and nature of these texts closely matches the type of training text data that we use to build the classifier (tweets). Note that none of the unlabeled tweets used for further training are present in the labeled training data. The amount of text needed to improve task-specific performance varies, but usually the more text, the more accurate the model becomes. Fourth, we save the new model and apply it to our classification task—building a classifier using labeled training data—in the same way we would apply original RoBERTa.[52]

We repeat this process for each of the Transformers models we compare in this article, BERT, RoBERTa, and DeBERTa, using each model's own tokenizer.[53] We compare the performance of the resulting three new models (BERT-Covid, RoBERTa-Covid, and DeBERTa-Covid) with that of the original models. We use a 10-times repeated 10-fold CV.[54]

The results are in table 3. We compare the performance of BERT, RoBERTa, and DeBERTa with (1) their respective Covid-specific further-trained models and (2) the same two baseline models from tables 1 and 2—SVMs and Bi-LSTM. The results confirm our expectations. First, all of the Transformers models outperform the baseline models significantly. BERT, the lowest performing Transformer model of the three, still outperforms SVM and Bi-LSTM-GloVe by 7% and 9.1%, respectively, in terms of their F1 scores when classifying fake news. Similar gains apply to true and undetermined news. RoBERTa and DeBERTa show even more significant gains in performance when compared to the two baselines.

The most relevant results are in comparing original BERT, RoBERTa, and DeBERTa with their respective Covid-trained models. All the Covid-trained models show substantively significant gains in performance when comparing F1 scores, especially when classifying fake news. First, BERT-Covid improves upon BERT by 4.22% in fake news, 1.81% in true news, and 3.45% in undetermined. RoBERTa-Covid bests RoBERTa by 3.91% in fake news, 3.45% in true news, and 3.37% in undetermined. Lastly, DeBERTa-Covid outperforms DeBERTa by 2.03% in fake news and by 2.56% in true news, and there is no

50. The dataset is available at https://github.com/MickeysClubhouse/COVID-19-rumor-dataset.

51. This choice is arbitrary though based on theory. Given the numerical representations of word embedding, we assume that virus + respiratory ≈ coronavirus. This is only a starting point for the new vocabulary elements. Once we further train the model, the embeddings for coronavirus and Covid will adjust and become more accurate representations of their meaning.

52. See appendix E for a more detailed step-by-step guide on how to further train a Transformers model with Python code snippets.

53. From the Transformers library, we use BertTokenizer for *bert–large–uncased*, RobertaTokenizerFast for *roberta–large*, and DebertaV2TokenizerFast for *deberta–v3–large*.

54. We use the same hyperparameters to train the original and new models to ensure the results are comparable.

improvement in undetermined news. Therefore, all further-trained models increase performance in classifying fake news around the Covid-19 pandemic by a range between 2.03% and 4.22%. RoBERTa sees the largest improvement in classifying true stories with 3.45%.

These results are especially significant in substantive terms for two reasons. First is that we trained the Transformers models with only 6,079 abstracts of academic articles on the topic of Covid-19, or 8.4 MB of text. We then added 4.8 million short news headlines in English that total 978 MB of text. Original RoBERTa, on the other hand, was trained on 160 GB of text. Our unlabeled data represent a fraction of the total training data in terms of quantity, and yet there are noticeable increases in overall performance. It is reasonable to expect larger performance gains with more training data or if we retrained the model from scratch. Second, our results are averages of 100 different models across 10 different sets of 10-fold CV. We can be certain that if we were to draw more model samples, the differences between the original and the Covid models would remain.

In all, these results show the potential that further training models holds for increasing performance in specialized applications. Differences across models are significant, especially considering that original BERT, RoBERTa, and DeBERTa already perform well in this application. Note, however, that performance will vary across domains, and researchers should decide whether further training the model is warranted for their specific application. We argue that doing so is important in domains where text requires complex contextual understanding or uses highly specialized language. Gordillo, Timoneda, and Vallejo Vera (2024), for instance, find that further training an XLM-RoBERTa model improved classifier performance by 8%. In other situations, in which researchers deem the performance of the original Transformers models to be sufficient, continuing without additional training can be an optimal choice.

In light of these results, our recommendation in this application is to fine-tune a RoBERTa model, further training it with unlabeled Covid-related data. To illustrate how to incorporate machine-labeled data in downstream statistical analyses, we provide an example by analyzing the effect of the length of a tweet in words on whether the content is fake news. We draw a random set of 3,000 news tweets around the coronavirus pandemic from the CoAID dataset by Cui and Lee (2020), who code each tweet as fake or true. These tweets are neither in the unlabeled dataset for further training nor in the final labeled dataset for standard fine-tuning. A total of 1,500 tweets are used as human-coded, gold-standard data. Another set of 1,500 are used as unlabeled data, and we apply our RoBERTa-Covid model to predict the labels. The final dependent variable is whether a tweet is fake news or not. The independent variable is tweet length. Further details of these tests as well as the complete set of results are in appendix I. The code is available from GitHub (omitted).[55] We find that there can be substantial discrepancies in model coefficients when using gold-standard or machine-coded labels, which is indicative of measurement bias. Applying Egami et al.'s (2023) test significantly alleviates measurement bias, and we recommend that scholars incorporate it into their workflow in downstream analyses using labels produced with fine-tuned Transformers models.

LIMITATIONS

Despite the substantial gains in performance and the flexibility of Transformers models, they have limitations. First, performance gains depend on setting the right hyperparameters for each application. This requires the researcher to perform various CV test runs to determine the learning rate that maximizes performance, minimizing both underfitting and overfitting.[56] Second, these models will most often be used to create categorical variables for downstream analyses. Since there is uncertainty and error around the Transformers model's classified labels, this error will then bleed into the analyses. We recommend following Egami et al.'s (2023) method to detect and correct for measurement bias in downstream statistical analysis. That said, any coding procedure will suffer from some form of coding error. Hence, researchers should focus on training models that produce the highest levels of accuracy for the categories that will then be used in their analyses. It is especially important to report variation across CV runs to understand the extent of the uncertainty around model results.

Third, the labeled training data in the final standard fine-tuning step that produces the classifier needs to be balanced for best performance. This may run counter to the balance of the categories in the real world, where some categories may be rare events but require a larger share of the training data for the model to operate correctly.[57] Fourth, NLP models are known to

55. The original Egami et al. (2023) code is available at osf.io/gjt87/?view_only=8f755cdf147f452a94297973eb83d85d.

56. Usually, a model overfits when the training loss is much smaller than the test loss. This means the model learns patterns in the training data too closely, thus minimizing training loss; however, those patterns may not exist in the test data, and therefore the model does not generalize well. Conversely, a model underfits when the training loss is larger than the test loss, meaning there is still more that the model can learn from the training data to generalize better to unseen data. CV helps find the right training balance that avoids both underfitting and overfitting.

57. Note that the training data need not be perfectly balanced. Rather, researchers need to keep the ratio between 2:1 and 3:1. With larger ration, the model will tend to overpredict the more frequent category to the detriment of the more infrequent one.

suffer from multimodality, which occurs when the "function we are trying to optimize is not globally concave" (Roberts et al. 2016, 2). Transformers-based models might suffer from multimodality too, hence the importance of repeated CV as a strategy to validate model results. Our results show that performance is quite consistent across different CV runs. Researchers must set hyperparameters carefully, using CV to determine the best set for their particular application (see app. B).[58]

Lastly, the performance gains from Transformers models come at a computational cost. The sheer amount of parallel calculations involved in the self-attention mechanism requires one or multiple GPUs, which have thousands of small cores and can perform all these calculations orders of magnitude faster than CPUs.[59] For smaller RoBERTa models, scholars can use Google Colab to run our sample code. Alternatively, they can access other cloud-based solutions with top-end GPUs, whose cost is now small and accessible.[60]

CONCLUSION

This article introduces Transformers-based classification models for English and multilingual text. We compare the performance of BERT, RoBERTa, and DeBERTa with other current state-of-the-art models used in political science and find multiple advantages. First, their capacity to understand context is greater than non-Transformers-based models. Understanding context increases model performance when generalizing to unseen data. We find that RoBERTa and DeBERTa generally perform much better than BERT. Second, the multilingual variants of different Transformers models, in particular XLM-R, perform exceptionally well when the training data is in multiple languages. Multilingual availability is important in comparative politics and international relations, especially with low-resource languages.[61] Third, Transformers models are flexible and can be further trained to fit a specific task, allowing researchers to "customize" a model to increase performance.

As stated above, the aim of this article is to compare the performance of recently developed Transformers models and highlight the versatility of these methods in social science research.[62] Furthermore, there are other practical benefits of using Transformers models. Given the model's overall improvements in out-of-sample accuracy, scholars can use them more consistently to turn their projects into big data projects, using manually labeled data as training data instead of as the full sample. This benefit is especially relevant with text data, which is often so abundant that the real limitations researchers face are related to resources and scaling costs. Transformers models—as other machine learning models—also allow researchers to understand their data much better.

We think the advantages of Transformers models greatly outweigh their limitations. This is especially true considering that the models introduced in this article represent a paradigm shift in NLP applications. They open up new avenues and opportunities for political science research across all subfields. Indeed, considering that text data is and will continue to be one of the biggest sources of data in the discipline, harnessing the power of these new models—and the ones that will inevitably follow—can have a transformational effect in applied political science research.

ACKNOWLEDGMENTS

The authors want to thank Logan Strother, Christina Walker, Ernesto Calvo, the participants at Purdue University's PSRW workshop series, and four anonymous reviewers for their helpful comments and suggestions. This work was completed in part with resources provided by Purdue University's Political Science Department/Rosen Center for Advanced Computing and the Research Computing Data Core at the University of Houston.

REFERENCES

Abercrombie, Gavin, and Riza Theresa Batista-Navarro. 2022. "Policy-Focused Stance Detection in Parliamentary Debate Speeches." In Leon Derczynski, ed., *Northern European Journal of Language Technology, V.8*. Copenhagen: Northern European Association of Language Technology.

58. In appendix B, we document the hyperparameters used to train our cases as reference for readers.

59. We understand that there is unequal access to computational resources in academia. While the costs of using Transformers-based models are not prohibitive (see app. A), they still pose a major obstacle to their use for research. Future work should aim to democratize these resources' use by lowering their cost.

60. See appendix D for details on available cloud-based solutions.

61. Transformers models allow political science researchers to overcome structural limitations of NLP when analyzing corpora of low-resource language—that is, languages lacking large monolingual or parallel corpora or manually crafted linguistic resources sufficient for building statistical NLP applications (Magueresse, Carles, and Heetderks 2020). The multilingual capabilities of models like mBERT and XLM-RoBERTa reduce the cost researchers have to incur in, for example, building new dictionaries or training Word2Vec-type embeddings.

62. In appendix G, we also showcase the application of our models in an established case by replicating Abercrombie's and Batista-Navarro's (2022) classification task of policy preferences in parliamentary speeches. By applying RoBERTa to their training set, we are able to increase performance considerably. The authors classify 34 topics to estimate policy preferences in speeches from the English Parliament. They apply a BERT-based model and obtain an F1 score of 50.9. We replicate their findings and apply a RoBERTa model, obtaining a 9.3 percentage point increase in the benchmark test. See appendix G for the performance table and an explanation of the replication process.

Abercrombie, Gavin, Federico Nanni, Riza Batista-Navarro, and Simone Paolo Ponzetto. 2019. "Policy Preference Detection in Parliamentary Debate Motions." In Mohit Bansal and Aline Villavicencio, eds., *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*. Hong Kong: Association for Computational Linguistics.

Albawi, Saad, Tareq Abed Mohammed, and Saad Al-Zawi. 2017. "Understanding of a Convolutional Neural Network." In *2017 International Conference on Engineering (ICET)*. Antalya, Turkey: IEEE, 1–6.

Alemán, Eduardo, Juan Pablo Micozzi Sebastián Vallejo Vera. 2022. "Congressional Committees, Electoral Connections, and Legislative Speech." *Political Research Quarterly* 76 (2): 994–1011.

Barberá, Pablo, Amber E. Boydstun, Suzanna Linn, Ryan McMahon, and Jonathan Nagler. 2021. "Automated Text Classification of News Articles." *Political Analysis* 29 (1): 19–42.

Beieler, John, Patrick T. Brandt, Andrew Halterman, Philip A. Schrodt, and Erin M. Simpson, 2016. "Generating Political Event Data in Near Real Time." In Michael Alvarez, ed., *Computational Social Science*. Cambridge: Cambridge University Press, 98.

Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. "Enriching Word Vectors with Subword Information." *Transactions of the Association for Computational Linguistics* 5:135–46.

Bonikowski, Bart, Yuchen Luo, and Oscar Stuhler. 2022. "Politics as Usual? Measuring Populism, Nationalism, and Authoritarianism in US Presidential Campaigns (1952–2020) with Neural Language Models." *Sociological Methods & Research* 51 (4): 1721–87.

Calvo, Ernesto, and Tiago Ventura. 2021. "Will I get COVID-19? Partisanship, Social Media Frames, and Perceptions of Health Risk in Brazil" *Latin American Politics and Society* 63 (1): 1–26.

Catalinac, Amy. 2016. *Electoral Reform and National Security in Japan: From Pork to Foreign Policy*. Cambridge: Cambridge University Press.

Chang, Charles, and Michael Masterson. 2020. "Using Word Order in Political Text Classification with Long Short-term Memory Models." *Political Analysis* 28 (3): 395–411.

Cheng, Mingxi, Songli Wang, Xiaofeng Yan, Tianqi Yang, Wenshuo Wang, Zehao Huang, Xiongye Xiao, Shahin Nazarian, and Paul Bogdan. 2021. "A COVID-19 Rumor Dataset." *Frontiers in Psychology* 12. https://doi.org/10.3389/fpsyg.2021.644801.

Conneau, Alexis, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. "Unsupervised Cross-lingual Representation Learning at Scale." Preprint, arXiv, https://doi.org/10.48550/arXiv.1911.02116.

Cui, Limeng, and Dongwon Lee. 2020. "CoAID: Covid-19 Healthcare Misinformation Dataset." Preprint, arXiv, https://doi.org/10.48550/arXiv.2006.00885.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. "BERT: Pre-training of Bidirectional Transformers for Language Understanding." Preprint, arXiv. https://doi.org/10.48550/arXiv.1810.04805.

Egami, Naoki, Musashi Jacobs-Harukawa, Brandon M. Stewart, and Hanying Wei. 2023. "Using Large Language Model Annotations for Valid Downstream Statistical Inference in Social Science: Design-Based Semi-Supervised Learning." Preprint, arXiv, https://doi.org/10.48550/arXiv.2306.04746.

Gervais, Bryan T., and Irwin L. Morris. 2019. "When Parties Tweet: Assessing and Explaining Incivility in State Legislatures." Presented at the Southern Political Science Association annual meeting, Austin, TX.

Gleason, Shane A. 2020. "Beyond Mere Presence: Gender Norms in Oral Arguments at the US Supreme Court." *Political Research Quarterly* 73 (3): 596–608.

González-Bailón, Sandra, and Georgios Paltoglou. 2015. "Signals of Public Opinion in Online Communication: A Comparison of Methods and Data Sources." *The ANNALS of the American Academy of Political and Social Science* 659 (1): 95–107.

Gordillo, Diana Dávila, Joan C. Timoneda, and Sebastián Vallejo Vera. 2024. "Machines Do See Color: A Guideline to Classify Different Forms of Racist Discourse in Large Corpora." Preprint, arXiv, https://doi.org/10.48550/arXiv.2401.09333.

Greer, Scott L., Elizabeth King, Elize Massard da Fonseca, and André Peralta-Santos. 2021. *Coronavirus Politics: The Comparative Politics and Policy of COVID-19*. Ann Arbor: University of Michigan Press.

Grimmer, Justin. 2010. "A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases." *Political Analysis* 18 (1): 1–35.

Grimmer, Justin, Margaret E. Roberts, and Brandon M. Stewart. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.

Hawkins, Kirk A., Rosario Aguilar, Bruno Castanho Silva, Erin K. Jenne, Bojana Kocijan, and Cristóbal Rovira Kaltwasser. 2019. "Measuring Populist Discourse: The Global Populism Database." Presented at EPSA annual conference, Belfast, UK.

He, Pengcheng, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. "DeBERTa: Decoding-Enhanced BERT with Disentangled Attention." Preprint, arXiv, https://doi.org/10.48550/arXiv.2006.03654.

Hobbs, William, and Nazita Lajevardi. 2019. "Effects of Divisive Political Campaigns on the Day-to-Day Segregation of Arab and Muslim Americans." *American Political Science Review* 113 (1): 270–6.

Johnson, Ben, Logan Strother, and Joan C. Timoneda. 2024. "Taking Justices' Legal Preferences Seriously." Preprint, arXiv (forthcoming).

Kaneko, Masahiro, and Danushka Bollegala. 2022. "Unmasking the Mask–Evaluating Social Biases in Masked Language Models." *Proceedings of the AAAI Conference on Artificial Intelligence* 36 (11): 11954–62.

Khan, Junaed Younus, Md Tawkat Islam Khondaker, Sadia Afroz, Gias Uddin, and Anindya Iqbal. 2021. "A Benchmark Study of Machine Learning Models for Online Fake News Detection." *Machine Learning with Applications* 4. https://doi.org/10.1016/j.mlwa.2021.100032.

King, Gary, Jennifer Pan, and Margaret Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 326–43.

Lazer, David MJ, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. 2018. "The Science of Fake News." *Science* 359 (6380): 1094–6.

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. "RoBERTa: A Robustly Optimized BERT Pretraining Approach." Preprint, arXiv, https://doi.org/10.48550/arXiv.1907.11692.

Loper, Edward, and Steve Bird. 2002. "NLTK: The Natural Language Toolkit." Preprint, arXiv, https://doi.org/10.48550/arXiv.cs/0205028.

Loshchilov, Ilya, and Frank Hutter. 2017. "Decoupled Weight Decay Regularization." Preprint, arXiv, https://doi.org/10.48550/arXiv.1711.05101.

Magueresse, Alexandre, Vincent Carles, and Evan Heetderks. 2020. "Low-Resource Languages: A Review of Past Work and Future Challenges." Preprint, arXiv, https://doi.org/10.48550/arXiv.2006.07264.

Matalon, Yogev, Ofir Magdaci, Adam Almozlino, and Dan Yamin. 2021. "Using Sentiment Analysis to Predict Opinion Inversion in Tweets of Political Communication." *Scientific Reports* 11 (1): 1–9.

McGregor, Shannon C. 2020. "Taking the Temperature of the Room: How Political Campaigns Use Social Media to Understand and Represent Public Opinion." *Public Opinion Quarterly* 84 (1): 236–56.

Medzihorsky, Juraj, Levente Littvay, and Erin K. Jenne. 2014. "Has the Tea Party Era Radicalized the Republican Party? Evidence from Text Analysis

of the 2008 and 2012 Republican Primary Debates." *PS: Political Science & Politics* 47 (4): 806–12.

Mikolov, Tomas, Edouard Grave, Piotr Bojanowski, Christian Puhrsch, and Armand Joulin. 2017. "Advances in Pre-training Distributed Word Representations." Preprint, arXiv, https://doi.org/10.48550/arXiv.1712.09405.

Pan, Jennifer, and Kaiping Chen. 2018. "Concealing Corruption: How Chinese Officials Distort Upward Reporting of Online Grievances." *American Political Science Review* 112 (3): 602–20.

Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. "GloVe: Global Vectors for Word Representation." In Alessandro Moschitti, Bo Pang, and Walter Daelemans, eds., *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computation Linguistics, 1532–43.

Ravichandiran, Sudharsan. 2021. *Getting Started with Google BERT*. Birmingham, UK: Packt Publishing.

Roberts, Margaret E., Brandon M. Stewart, and Dustin Tingley. 2016. "Navigating the Local Modes of Big Data: The Case of Topic Models." In R. Michael Alvarez, ed., *Computational Social Science*. Cambridge: Cambridge University Press, 51–97.

Rodriguez, Pedro L., and Arthur Spirling. 2022. "Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research." *The Journal of Politics* 84 (1): 101–15.

Strother, Logan. 2017. "How Expected Political and Legal Impact Drive Media Coverage of Supreme Court Cases." *Political Communication* 34 (4): 571–89.

Timoneda, Joan C., and Sebastián Vallejo Vera. 2021. "Will I Die of Coronavirus? Google Trends Data Reveal that Politics Determine Virus Fears." *PLoS ONE* 16 (10): e0258189.

Tunstall, Lewis, Leandro von Werra, and Thomas Wolf. 2022. *Natural Language Processing with Transformers*. Sebastopol, CA: O'Reilly Media, Inc.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. "Attention is All You Need." In Ulrike von Luxburg and Isabelle Guyon, eds., *Proceedings of the 31st International Conference on Neural Information Processing Systems*. Red Hook, NY: Curran Associates Inc., 6000–10.

Ward, Michael D., Andreas Beger, Josh Cutler, Matthew Dickenson, Cassy Dorff, and Ben Radford. 2013. "Comparing GDELT and ICEWS Event Data." *Analysis* 21 (1): 267–97.

Wettig, Alexander, Tianyu Gao, Zexuan Zhong, and Danqi Chen. 2022. "Should You Mask 15% in Masked Language Modeling?" Preprint, arXiv, https://doi.org/10.48550/arXiv.2202.08005.

Zhang, Tianyi, Tatsunori Hashimoto. 2021. "On the Inductive Bias of Masked Language Modeling: From Statistical to Syntactic Dependencies." Preprint, arXiv, https://doi.org/10.48550/arXiv.2104.05694.